

10/PRFS

PROCESS ORIENTED COMPUTING ENVIRONMENT

This invention relates to a computing environment, and in particular to a process oriented computing environment, and also to a method for controlling the migration and evolution of processes within the environment.

Recent years have seen a number of developments in computing science regarding how elements within a software application are treated and handled. In this context the most basic elements to be found within a software application are data and program modules. Traditional procedural programming paradigms focus on the logic of the software application, so a program is structured based on program modules. One uses the program modules by explicitly supplying to them the data on which they should operate. A potential pitfall of this paradigm is that it is difficult to guarantee type safety, ie that the data being passed to a program module is of the correct type.

More recently there has been a move towards an object-oriented paradigm. In this paradigm, programs are structured around objects, with each object representing an entity in the world being modeled by the software application. Each object manages its own data (state), which are hidden from the external world (other objects and programs). An object in a program interacts with another object by sending it a message to invoke one of its exposed program modules (method). This paradigm imposes better control and protection over internal data, and helps to structure complex applications designed and implemented by a team of programmers. An example of an object-oriented environment can be found in US 5,603,031. This discloses an environment in which new agents (essentially objects) consisting of data and program modules can be sent between machines.

While object-oriented paradigm represents a significant advance in software engineering, the data and modules that constitute each object are static. The paradigm is still inadequate for writing programs that must evolve during execution, eg programs that need to pick up, drop, or substitute selected modules. There have been several attempts at overcoming this limitation. For example, work described in US Patents 4954941, 5175828, 5339430 and 5659751 address techniques for re-linking or re-binding selected software modules dynamically during runtime. Also Microsoft's Win32 provides for

explicit mapping and un-mapping of dynamic linked libraries into the address space of a process through the LoadLibrary and FreeLibrary calls. With this prior art, however, the prototype or specification of functions and symbols are fixed beforehand and compiled into application programs. This means that an object cannot invoke a module of another object for which the specification is not known at compile time.

Another shortcoming of both traditional procedural and object-oriented paradigms is that programs consist only of data and program modules, but not the transient information that capture the entire state of execution during runtime. This makes it difficult to interrupt a running program and to migrate it to another machine. Generally it is only possible to transfer the saved data file of completed applications. As a very simple example of this problem, while a word processing document file or a spreadsheet file may be transferred from one machine to another, it is not possible to do so while the file is currently being worked on without reverting to a saved version. Transient information is lost. In the case of a word processing file, for example, this means that any "undo editing" function cannot be used by the receiving machine on the file it has just received because the transient "undo" information is not part of the data file.

At present, such migrations have been effected from outside the program, as utilities in the computing environment. Examples include Amoeba, Charlotte, Sprite and Condor. Such utilities work by taking a snapshot of the running program (or core dump), and by resuming execution on the recipient machine from the snapshot. Unfortunately the migration utilities have no knowledge of the usage requirements and semantics of the components of the running program, and so there is no way to adapt it to the new computing environment. Consequently, the migrated program most likely cannot run in a different computing environment from the original one, such as when the machines have different devices like displays, hard disks and sound cards. Even in cases where it does, it would not run with the same level of efficiency, for example because the machines have different amounts of main memory.

In this specification the following terms will be used with the following meaning:

"First class entity": an object that can be manipulated directly.

"Process": a combination of data, program module(s) and current execution state.

"Execution state": the values and contents of transient parameters such as the contents of registers, frames, counters, look-up tables and the like.

It is an object of the present invention to provide a computing environment that treats data, program modules and processes all as first class entities and which facilitates the migration and evolution of a process during runtime.

According to the present invention there is provided a computing environment wherein a process comprising a set of data and/or program modules and execution state is treated as an entity that can be transferred between components of the environment, and wherein a said process can be subject to evolutionary operations.

By means of this arrangement a process itself is treated as a first class entity in addition to the traditional first class entities of data and program modules, and the first class entities of object-oriented programming that combine data and program modules. Thus a process can be transferred between machines together with all transient runtime data that comprises the current execution state. This allows an application to be transferred between machines during runtime and to continue running without any interruptions. In addition, however, not only can a process migrate but it can also be subject to evolutionary operations and can thus evolve, either within a first machine or during migration. In addition not only can a process be transferred between different hardware components directly, but a process may be transferred to a memory device for storage preceding subsequent transfer to another hardware component or even back to the first component. Furthermore processes or sub-processes forming a subset of a process may be transferred between processes themselves, and such processes may be running on different hardware components or may be all running within a single hardware component. It is important to appreciate therefore that in the present invention the components of the environment between which a process may be transferred may comprise hardware and/or software elements of the computing environment.

In a preferred embodiment of the invention a construct is formed comprising a set of data and/or program modules and execution state of a first process, and wherein the evolutionary operations are performed by functions operating on a said construct. In particular the construct is formed by a construct operation that suspends all active threads of the first process and creates a new process comprising at least some of the data and/or

program modules and execution state of the first process, and stores the new process in a data area of the said process. While the construct may comprise all of the first process, more commonly the construct will comprise only data, program modules and execution state falling within lists that are passed to the construct operation. Preferably the construct may also be formed with an authorising signature in order to confirm its authenticity if it is to be transferred to another component of the computing environment.

Simple evolutionary operations can be performed on the process stored in the construct without it being transferred to any other component. In particular the process can be modified by either selective deletion of objects from within the process, and/or by selective loading or reloading of object into the process.

In more sophisticated embodiments of the invention, however, a process may be subject to evolutionary operations that include the incorporation into a first process of objects from a second process. Again this may be achieved by firstly forming a construct comprising at least some of the data and/or program modules and execution state from the second process and then transferring this construct to the first process. Preferably the construct is formed by a construct operation that suspends all active threads of said second process and creates a new process comprising some of the data and/or program modules and execution state of said second process, and stores said new process in a data area of said second process. Again while the construct may comprise all of the second process, it may also comprise only data, program modules and execution state falling within lists that are passed to said construct operation. An authorising signature may also be provided to the construct.

There are, however, a number of ways that this can be done and in which the elements of the two processes can be combined.

In a simple embodiment after the construct is transferred the second process stored within the construct is caused to be activated within said first process.

In another embodiment, however, after the construct is transferred the first process is suspended and the second process stored within the construct is activated, and when the second process is concluded the data and program modules of the second process are added to the first process and the first process is re-activated.

0055513 061504
T05750 FTS5500

Another possibility is that after the first process terminates, a subset of data and program modules from the first process are added to the second process stored in the construct and the second process is then activated.

It may not always be necessary to transfer information relating to execution state. Therefore in another embodiment a construct is formed comprising data and/or program modules from the second process, and the construct is transferred to the first process. In this embodiment only data and program modules are transferred between processes. As in other embodiments the construct is formed by a construct operation that suspends all active threads of the second process and creates a new process comprising at least some of the data and program modules of the second process, and stores the new process in a data area of the second process. Again the construct may comprise all of the data and program modules from the second process, or may comprise only a subset of the data and program modules in which case the construct comprises only data and program modules falling within lists that are passed to the construct operation. In this embodiment the data and program modules from the second process may be copied into the first process.

When data, program modules and execution state information is being interchanged between different processes it is possible that there may be conflicts, for example between data values of two processes. Therefore preferably flagging techniques may be employed such that in the event of a conflict one process is given priority over the other such that its data, program modules and/or execution state will override those of the other process in the event of a conflict.

It should also be understood that the operations described herein are complementary and may be combined in many different ways. For example, a first process may modify itself by deletion of unwanted program modules, followed by the loading of new program modules and data, and then the resulting new process may be transferred to a second process within which it is assimilated. Equally one process may be transferred to another process and combined with it to form a new process, which then be transferred to combine with a third process and so on.

Viewed from another broad aspect it will be seen that the invention also provides a method for controlling the evolution and migration of a process comprising a set of data, program modules and execution state within a computing environment, wherein a

construct is formed comprising at least some of the data and/or program modules and execution state of a first process.

Thus it will be seen that the present invention provides techniques for the controlled evolution and migration of processes within a computing environment, which environment may be a distributed computing environment. The present invention therefore has a wide range of potential applications, including for example the creation of software applications having far greater flexibility than is known with current techniques. For example, the capability of a process is no longer fixed at compilation, but rather a process can evolve by acquiring new functionality and/or by shedding unwanted elements. Instead of transferring only data and/or program modules as with conventional object-oriented technology, a process including execution state can be transferred and new processes can be created in a semi-executed state which allows the past history of a process to be propagated. Such a technique has a number of advantages in allowing, for example, the creation of a genuine "plug-and-play" environment, and also the customisation of software in realtime without requiring recompilation or relinking.

Some embodiments of the present invention will now be described by way of example and with reference to the accompanying drawings, in which:

Fig.1 is a general schematic model of an operating environment of a computing system,

Fig.2 schematically illustrates a process life-cycle and operations that may be performed on a process,

Fig.3 is a flow-chart illustrating the operation Hibernaculum Construct(Stack s, Module m, Data d),

Fig.4 is a flow-chart illustrating operation int Assimilate(Hibernaculum h, Override Flags f),

Fig.5 is a flow-chart illustrating the operation int Usurp(Hibernaculum h, Override Flags f),

Fig.6 is a flow-chart illustrating the operation int Bequeath(Hibernaculum h, Thread threadname, Override Flags f),

Fig.7 is a flow-chart illustrating the operation int Inherit(Hibernaculum h, Thread threadname, Override Flags f),

Fig.8 is a flow-chart illustrating the operation `int Mutate(Stacks, int sFlag, Module m, int mFlag, Data d, int dFlag)`,

Fig.9 is a flow-chart illustrating the operation `int Checkpoint(Process p1, Target t)`, and

5 Fig.10 is a flow-chart illustrating the operation `int Migrate(Process p1, Machine m)`.

Figure 1 shows the general model of a computing system. An application program
30 comprises data 10 and program modules 20. The operating system 60, also known as
10 the virtual machine, executes the application 30 by carrying out the instructions in the
program modules 20, which might cause the data 10 to be changed. The execution is
effected by controlling the hardware of the underlying machine 70. The status of the
execution, together with the data and results that the operating system 60 maintains for
the application 30, form its execution state 40.

15 Such a model is general to any computing system. It should be noted here that the
present invention starts from the realisation that all the information pertaining to the
application at any time is completely captured by the data 10, program modules 20 and
execution state 40, known collectively as the process 50 of the application 30.

20 The process 50 can have one or more threads of execution at the same time. Each
thread executes the code of a single program module at any given time. Associated with
the thread is a current context frame, which includes the following components:

- A set of registers
- A program counter, which contains the address of the next instruction to be executed
- 25 • Local variables of the module
- Input and output parameters of the module
- Temporary results of the module

30 In any module A, the thread could encounter an instruction to invoke another module
B. In response, the program counter in the current frame is incremented, then a new
context frame is created for the thread before it switches to executing module B. Upon

completing module B, the new context frame is discarded. Following that, the thread reverts to the previous frame, and resumes execution of the original module A at the instruction indicated by the program counter, i.e., the instruction immediately after the module invocation. Since module B could invoke another module, which in turn could invoke some other module and so on, the number of frames belonging to a thread may grow and reduce with module invocations and completions. However, the current frame of a thread at any given time is always the one that was created last. For this reason, the context frames of a thread are typically stored in a stack with new frames being pushed on and popped from the top. The context frames of a thread form its execution state, and the state of all the threads within the process 50 constitute its execution state 40 in Fig. 1.

The data 10 and program modules 20 are shared among all threads. The data area is preferably implemented as a heap, though this is not essential. The locations of the data 10 and program modules 20 are summarized in a symbol table. Each entry in the table gives the name of a datum or a program module, its starting location in the address space, its size, and possibly other descriptors. Instead of having a single symbol table, each process may alternatively maintain two symbol tables, one for data alone and the other for program modules only, or the process could maintain no symbol table at all.

In a preferred embodiment of the present invention, the data and program code of a process are stored in a heap and a program area respectively and are shared by all the threads within the process. In addition the execution state of the process comprises a stack for each thread, each stack holding context frames, in turn each frame containing the registers, local variables and temporary results of a program module, as well as addresses for further module invocations and returns. Before describing an embodiment of the invention in more detail, however, it is first necessary to introduce some definitions of data types and functions that are used in the embodiment and which will be referred to further below.

In addition to conventional data types such as integers and pointer, four new data types Data, Module, Stack and Hibernaculum are defined in the present invention:

30 Data: A variable of this data type holds a set of data references. Members are added to and removed from the set by means of the following functions;

Int AddDatum(Data d, String dataname) inserts the data item dataname in the heap of the process as a member of d.

Int DelDatum(data d, String dataname) removes the data item dataname from d.

- 5 Module: A variable of this data type holds a set of references to program modules. Members are added to and removed from the set with the following functions;

Int AddModule(Module d, String modulename) inserts the program module modulename in the program area of the process as a member of d.

10

Int DelModule(Module d, Stringname modulename) removes the program module modulename from d.

15

Stack: A variable of this data type holds a list of ranges of execution frames from the stack of the threads. The list may contain frame ranges from multiple threads, however no thread can have more than one range. Variables of this type are manipulated by the following functions:

20

Int OpenFrame(Stack d, Thread threadname) inserts into d a new range for the thread threadname, beginning with the thread's current execution frame. This function has no effect if the thread already has a range in d.

25

Int CloseFrame(Stack d, Thread threadname) ends the open-ended range in d that belongs to the thread threadname. This function has no effect if the thread does not currently have an open-ended range in d.

Int PopRange(Stack d, Thread threadname) removes from d the range belonging to the thread threadname.

- 30 Hibernaculum: A variable of this data type is used to hold a suspended process.

00355513-051504
T05T90-ET55300

10

- 10

10

- 15

20

20

- 25

- 30

environment, the override flags specify which to preserve. Fig.5 is a flow-chart illustrating the steps of the usurp operation.

Int Bequeath(Hibernaculum h, Thread threadname, OverrideFlags f), upon threadname's termination, activates the threads of the process stored within h and runs them as threads within the environment of the process containing threadname. Where there is a conflict between the data and/or program modules of the hibernaculum and the calling process, the override flags specify which is to be preserved. Fig.6 is a flow-chart illustrating the steps of the bequeath operation.

Int Inherit(Hibernaculum h, Thread threadname, OverrideFlags f) suspends threadname and activates the process within h. When the process within h terminates its data and program modules are added to the process containing threadname before threadname is reactivated. Where there is a conflict between the data and/or program modules of the hibernaculum and the process containing threadname, the override flags specify which is to be preserved. Fig.7 is a flow-chart illustrating the steps of the inherit operation.

Int Mutate(Stack s, int sFlag, Module m, int mflag, Data d, int dflag) modifies the execution state, program table and data heap of the calling process. If a thread has an entry in s, only the range of execution frames specified by this entry is preserved, the other frames are discarded. Execution stacks belonging to threads without an entry in s are left untouched. In addition, program modules listed in m and data items listed in d are kept or discarded depending on the flag status. Fig.8 is a flow-chart illustrating the steps of the mutate operation.

Int Checkpoint(Process p, Target t) creates a snapshot of p including all of its data, program modules and execution state. The snapshot is then sent to the specified target.

Int Migrate(Process p, Machine m) transfers the process p to the specified machine for p to continue execution there. All of the data (including file handles and established sockets), program modules and execution state are preserved in the transfer.

The typical life cycle of a process in an embodiment of the present invention is depicted in Figure 2. First, an application 210 or a suspended process 220 stored in a hibernaculum is loaded by the operating system, then starts running as a process 230. The application 210 may be a program that is designed to perform some tasks, or it may be a simple loader that is used to start up other processes, while the suspended process 220 could be either produced locally earlier or generated on and transferred from another machine. As the process 230 goes through the application program logic, the process may create new processes, absorb other processes, mutate, or migrate as explained below.

In particular, the process may be stored in a construct hibernaculum in a suspended state. In this condition the process may be subject to a number of operations. To begin with the hibernaculum may be sent to another operating environment. Alternatively the process may be modified within its own operating environment by the selective deletion of elements from within the process, or by the selective reloading of elements into the process. A still further possibility is that the process may receive a suspended process from another operating environment, and either all or part of this suspended process may be incorporated into the first process. Alternatively all or part of the first process may be incorporated into the second process. It will also be understood that all these operations may be combined in many different ways. For example, a process may be sent from one operating environment to another and then may mutate by dropping certain elements and reloading other elements when in the second environment. In the following description the construction of a hibernaculum, send and receive functions, and exemplary evolutionary operations will all be described in greater detail.

New processes are created with the Construct 110 operation. Fig.3 is a flow-chart showing the steps of this Construct operation. Each invocation of this operation starts up a controller thread in the process 230. The controller thread freezes all other active threads in the process 230, then creates a new process with some or all of the execution state, program modules and/or data of the process 230 except for those belonging to the controller thread, before resuming the frozen threads. Therefore, the new process contains no trace of the controller thread. The new process is suspended immediately and returned in a hibernaculum in the data area of the process 230. As explained earlier, a

hibernaculum is a special data type that serves the sole purpose of providing a container for a suspended process. Since a process may have several hibernacula in its data area, it could create a new hibernaculum that contains those hibernacula, each of which in turn could contain more hibernacula, and so on. When the new process is activated subsequently, only those threads that were active just before the Construct 110 operation will begin to execute initially; threads that were suspended at that time will remain suspended. At the end of the Construct 110 operation, the controller thread resumes those threads that were frozen by it before terminating itself.

To specify what execution state should go into the new process, the Construct 110 operation is passed a list of ranges of context frames. The list may include frames from the state of several threads. No thread is allowed to have more than one range in this list. A thread can specify that all of its frames at the time of the Construct 110 operation are to be included in the list, by calling the AllFrame function beforehand. Alternatively, the thread can call the OpenFrame function to register its current frame, so that all the frames from the registered frame to the current frame at the time of the Construct 110 operation are included in the list. The thread can also call the CloseFrame function subsequently, to indicate that only frames from that registered with OpenFrame to the current frame at the time of calling CloseFrame are to be included in the list for the Construct 110 operation. An AllFrame or OpenFrame request for a thread erases any previous AllFrame, OpenFrame and CloseFrame requests for that thread. A CloseFrame request overrides any earlier CloseFrame request for the same thread, but the request is invalid if the thread has not already had an OpenFrame request. A thread can also make AllFrame, OpenFrame and/or CloseFrame requests on behalf of another thread by providing the identity of that thread in the requests; the effect is as if that thread is making those requests itself.

The Construct 110 operation can also be passed a list of program modules that should go into the newly created process. A thread can specify that all modules of the process are to be included in the list, by calling the AllModules function prior to the Construct 110 operation. Alternatively, the thread can call the AddModule function to add a specific module to the list, and the DelModule function to remove a specific module from the list. The effect of the AllModules, AddModule and DelModule requests,

possibly made by different threads, are cumulative. Hence a DelModule request after an AllModules request would leave every module in the list except for the one removed explicitly, and a DelModule can be negated with an AddModule or AllModules request. As there could be multiple AddModule requests for the same module and AllModules could be called multiple times, a program module may be referenced several times in the list. However, the Construct 110 operation consolidates the entries in the list, so no program module gets duplicated in the new process.

To copy some or all of the data of the process 230 to the new process, the Construct 110 operation can be passed a data list. This list contains only the name of, or reference to data that should be copied. The actual data content or values that get copied to the new process are taken at the time of the Construct 110 operation, not at the time that each datum is added to the list. To ensure consistency among data that could be related to each other, all the threads in the process 230 are frozen during the Construct 110 operation. A thread can specify that all data of the process 230 are to be included in the list, by calling the AllData function prior to the Construct 110 operation. Alternatively, the thread can call the AddDatum function to add a specific datum to the list, and the DelDatum function to remove a specific datum from the list. The effect of the AllData, AddDatum and DelDatum requests, possibly made by different threads, are cumulative. Hence a DelDatum request after an AllData request would leave all of the data in the list except for the one removed explicitly, and a DelDatum can be negated with an AddDatum or AllData request. As there could be multiple AddDatum requests for the same datum and AllData could be called multiple times, a datum may be referenced several times in the list. However, the Construct 110 operation consolidates the entries in the list, so no datum gets duplicated in the new process.

Since the lists passed to the Construct 110 operation are constructed from the execution state, program modules and data of the process 230, the new process initially does not contain any component that is not found in the process 230. Consequently, the symbol table in the new process is a subset of the symbol table of the process 230. Threads in the process 230 that do not have any frame in the new process are effectively dropped from it. For those threads that have frames in the new process, when activated later, each will begin execution at the instruction indicated by the program counter in the

most recent frame amongst its frames that are copied. By excluding one or more of the most recent frames from the new process, the associated thread can be forced to return from the most recent module invocations. An exception is raised to alert the thread that those modules are not completed normally. Alternatively, the thread can be made to redo those modules upon activation, by decrementing the program counter in the most recent frame amongst those frames belonging to that thread that are copied. Similarly, by excluding one or more of its oldest frames from the new process, a thread can be forced to terminate directly after completing the frames that are included.

Hibernacula, produced by the Construct 110 operation, can be exchanged between processes via the Send 120 and Receive 130 operations. The process 230 can send a hibernaculum in its data area out on an output stream, by invoking the Send 120 operation with the hibernaculum and output stream as parameters. The output stream could lead to either a disk file, a memory stream, a device stream, or the input stream of another process, possibly on a different machine. In the former case, the process in the hibernaculum is stored away for the time being, whereas in the latter case the process in the hibernaculum is received directly into the data area of another process.

The process 230 can also acquire other processes via the Receive 130 operation, which receives from an input stream a hibernaculum containing a suspended process. The source of the input stream could be a disk file, a memory stream, a device stream, or the output stream of another process, possibly from a different machine.

The process 230 can absorb the suspended process within a hibernaculum through the Assimilate 140, Usurp 150, Bequeath 160 and Inherit 170 operations. These operations can be invoked by any thread in the process 230, or by another process that has appropriate permissions. The hibernaculum could have been created by the process 230 itself earlier, or received from another process via a disk file or an input stream.

The Assimilate 140 operation accepts a hibernaculum as input. Fig.4 is a flow-chart showing this operation in detail. The operation starts up a controller thread in the process 230. The controller thread freezes all other active threads in the process 230, adds the program modules and data of the process in the hibernaculum to the program modules and data of the process 230, respectively, and updates its symbol table accordingly. In case of a name conflict, the program module or data from the hibernaculum is discarded

in favor of the one from the process 230 by default. However, a flag can be supplied to give preference to the hibernaculum. Moreover, the context frames of the threads within the hibernaculum are added to the execution state of the process 230, enabling those threads to run within the process. Only those threads that were active just before the hibernaculum was created are activated initially; threads that were suspended at that time remain suspended. Each newly acquired thread begins execution at the instruction indicated by the program counter in the most recent frame amongst the context frames that belong to that thread. Finally, all the original threads in the process 230 that were frozen by the controller thread are resumed before it terminates itself.

The Usurp 150 operation, which also accepts a hibernaculum as input, enables the process 230 to take in only program modules and data from a hibernaculum, without acquiring its threads. Fig.5 is a flow-chart showing this operation in detail. The operation starts up a controller thread in the process 230. The controller thread freezes all other active threads in the process 230, adds the program modules and data of the process in the hibernaculum to the program modules and data of the process 230, respectively, and updates its symbol table accordingly. In case of a name conflict, the program module or data from the hibernaculum is discarded in favor of the one from the process 230 by default. However, a flag can be supplied to give preference to the hibernaculum. Finally, all the original threads in the process 230 that were frozen by the controller thread are resumed before it terminates itself.

The Bequeath 160 operation accepts a hibernaculum and a thread reference as input. Fig.6 is a flow-chart showing this operation in detail. It starts up a bequeath-thread in the process 230. The bequeath-thread registers the referenced thread and any existing bequeath-threads on that thread, then allows them to carry on execution without change. After all those threads and threads that they in turn activate have terminated, the bequeath-thread loads in the context frames, program modules and data in the hibernaculum. In case of a name conflict, the program module or data from the hibernaculum is discarded in favor of the one from the process 230 by default. However, a flag can be supplied to give preference to the hibernaculum. Subsequently, threads within the hibernaculum are activated to run in the process 230 before the bequeath-thread terminates itself. Only those threads that were active just before the hibernaculum

0055416104
T06T90:ET55580

was created are activated initially; threads that were suspended at that time remain suspended. Each newly acquired thread begins execution at the instruction indicated by the program counter in the most recent frame amongst the context frames that belong to that thread. If multiple Bequeath 160 requests are issued for the same thread, there will be several bequeath-threads in the process 230. Each bequeath-thread will wait for all the existing bequeath-threads on the same thread, together with all the threads that they activate, to terminate before performing its function. As a result, the Bequeath 160 requests are queued up and serviced in chronological order.

The reverse of Bequeath 160 is the Inherit 170 operation, which also accepts a hibernaculum and a thread reference as input. The flow-chart for this operation is shown in Fig.7. This operation starts up an inherit-thread in the process 230. The inherit-thread freezes the referenced thread in the process 230 together with the bequeath-threads and any inherit-thread for the referenced thread, adds the program modules and data in the hibernaculum to the program modules and data of the process 230, respectively, and updates its symbol table accordingly. In case of a name conflict, the program module or data from the hibernaculum is discarded in favor of the one from the process 230 by default. However, a flag can be supplied to give preference to the hibernaculum. Moreover, the context frames of the threads within the hibernaculum are added to the execution state of the process 230, enabling those threads to run within the process. Only those threads that were active just before the hibernaculum was created are activated initially; threads that were suspended at that time remain suspended. Each newly acquired thread begins execution at the instruction indicated by the program counter in the most recent frame amongst the context frames that belong to that thread. After all the acquired threads and threads that they in turn activate have terminated, the inherit-thread resumes those threads that were frozen by it earlier, before terminating itself. If one of the acquired threads issues an Inherit 170 request, the acquired threads and the current inherit-thread would in turn be suspended, pending the completion of the latest Inherit operation. If an acquired thread does a Bequeath 160, the inherit-thread would wait for the Bequeath request to be satisfied before resuming the threads that were frozen by it.

Besides constructing and absorbing new processes, the process 230 can also modify any of its components directly by calling the Mutate 180 operation from any thread. Fig.8

shows the flow-chart for the Mutate operation. The operation starts up a controller thread in the process 230. The controller thread first freezes all other active threads in the process 230, then selectively retains or discards its execution state, program modules and data. A list of context frames, together with a flag, can be passed to the Mutate 180 operation to retain or discard the frames in the list. Similarly, a program module list and/or a data list can be provided to indicate program modules and data of the process 230 that should be retained or discarded. The generation of the execution state, program module and data lists are the same as for the Construct 110 operation. After mutation, the controller thread resumes the threads that were frozen by it before terminating itself.

Threads that no longer have a context frame in the process 230 are terminated. Each of the remaining threads resumes execution at the instruction indicated by the program counter in the most recent frame amongst the retained frames belonging to that thread. By discarding one or more of its most recent frames, a thread can be forced to return from the most recent module invocations. An exception is raised to alert the thread that those modules are not completed normally. Similarly, by discarding one or more of its oldest frames, a thread can be forced to terminate directly after completing the frames that are retained. Space freed up from the discarded context frames, program modules and data is automatically reclaimed by a garbage collector.

The process 230 is also capable of suspending and migrating itself. The Checkpoint 190 operation starts up a controller thread in the process 230. Fig.9 is a flow-chart showing the Checkpoint operation in detail. The controller thread freezes all other active threads in the process 230, then sends it in the form of a hibernaculum on a specified output stream that it established earlier. The hibernaculum contains all of the execution state, program modules, and data of the process 230, except for those belonging to the controller thread and the output stream used by the Checkpoint 190 operation. The output stream could lead to either a disk file or the input stream of another process, possibly on a different machine. The controller thread ends by terminating the process 230. When the process in the hibernaculum is activated subsequently, only those threads that were active just before the Checkpoint 190 operation will begin to execute initially; threads that were suspended at that time will remain suspended. The Checkpoint 190 operation could be

5

10

15

20

To implement the process migration and adaptation system of the present invention in a Java environment, a package called snapshot is introduced. This package contains the following classes, each of which defines a data structure that is used in the migration and adaptation operations:

25

30

```
public class State {
    ...
}
```

5

...

}

10

...

}

10

...

}

15

20

20

20

20

}

25

25

25

25

25

25

25

30

30

```
public static native int Migrate(Machine m);
```

```
// This class is not to be instantiated
```

```
private Snapshot() {
```

```
5 }
```

```
}
```

The methods in the Snapshot class can be invoked from application code. For example:

```
10 try {
```

```
    if (snapshot.Snapshot.Construct(s, m, d) != null) {
```

```
        // hibernaculum has been created
```

```
    } else {
```

```
15        // failed to create hibernaculum
```

```
    }
```

```
    catch(snapshot.SnapshotException e) {
```

```
        // Failed to create hibernaculum
```

```
    }
```

The migration and adaptation operations are implemented as native codes that are added to the Java virtual machine itself, using the Java Native Interface (JNI). To do that, a Java-to-native table is first defined:

```
25 #define KSH "Ljava/snapshot/Hibernaculum;"
```

```
#define KSS "Ljava/snapshot/State;"
```

```
#define KSM "Ljava/snapshot/Module;"
```

```
#define KSD "Ljava/snapshot/Data;"
```

```
30 static JNINativeMethod snapshot_Snapshot_native_methods[] = {
```

```
{
```

0985513-061501
T05T90-ET5550

```
    "Construct",
    ("("KSSKSMKSD")"KSH,
    (void*)Impl_Snapshot_Construct
},
5    {
    "Send",
    ("("KSH"Ljava/io/OutputStream;)I",
    (void*)Impl_Snapshot_Send
},
10    {
    "Receive",
    ("(Ljava/io/InputStream;)"KSH,
    (void*)Impl_Snapshot_Receive
},
15    {
    "Assimilate",
    ("("KSH"I)I",
    (void*)Impl_Snapshot_Assimilate
},
20    {
    "Usurp",
    ("("KSH"I)I",
    (void*)Impl_Snapshot_Usurp
},
25    {
    "Bequeath",
    ("("KSH"I)I",
    (void*)Impl_Snapshot_Bequeath
},
30    {
    "Inherit",
```

0955543-051501
TOST90-ET55880

}

{

5

}

{

10

}

{

15

}

} .

20

JNIEXPORT void JNICALL

```
Java_snapshot_Snapshot_registerNatives(JNIEnv *env, jclass cls) {
```

25

30

Besides the above native codes, several functions are added to the Java virtual machine implementation, each of which realizes one of the migration and adaptation operations:

```
5  void* Impl_Snapshot_Construct(..) {  
    // follow flowchart in Figure 3  
    ...  
}  
  
10 void* Impl_Snapshot_Send(..) {  
    // send given hibernaculum to specified target  
    ...  
}  
  
15 void* Impl_Snapshot_Receive(..) {  
    // receive a hibernaculum from a specified source  
    ...  
}  
  
20 void* Impl_Snapshot_Assimilate(..) {  
    // follow flowchart in Figure 4  
    ...  
}  
  
25 void* Impl_Snapshot_Usurp(..) {  
    // follow flowchart in Figure 5  
    ...  
}  
  
30 void* Impl_Snapshot_Bequeath(..) {  
    // follow flowchart in Figure 6
```



```
...  
}
```

```
void* Impl_Snapshot_Inherit(..) {  
5      // follow flowchart in Figure 7  
      ...  
}
```

```
void* Impl_Snapshot_Mutate(..) {  
10     // follow flowchart in Figure 8  
      ...  
}
```

```
void* Impl_Snapshot_Checkpoint(..) {  
15     // follow flowchart in Figure 9  
      ...  
}
```

```
void* Impl_Snapshot_Migrate(..) {  
20     // follow flowchart in Figure 10  
      ...  
}
```

25

30

0055513-061501